

Efficient Smoothing of Unstructured Grids

Markus Berndt and J. David Moulton
(T-7); moulton@lanl.gov

The smoothing of unstructured grids in large-scale 2D and 3D simulations is a critical component of many projects related to the Advanced Simulation and Computing (ASC) Program. Many techniques have been developed to perform this smoothing, but few are capable of handling large unstructured meshes in complex geometries. For example, most smoothers generate unacceptable or even invalid grids in the neighborhood of extremely convex or concave boundaries. In contrast, a smoothing technique was developed [1, 2] that is robust with respect to these geometric complexities. This grid smoothing methodology is based on the concept of harmonic coordinates, and hence, has a natural variational formulation. Specifically, this formulation defines a system of quasilinear elliptic partial differential equations (PDEs) with one equation for each of the grid coordinates. The only coupling in this system arises through the components of the metric tensor, which plays the role of a diffusion tensor. Thus, the system may be described as a set of quasilinear diffusion equations with a solution-dependent diffusion tensor. The grid smoothing is driven by defining a target metric, which in 2D is based on a dual grid. Since the quasilinear system of PDEs arises naturally from a variational formulation, a standard vertex-based Finite Element Method (FEM) is used to discretize this nonlinear system.

The Newton-Krylov approach is used to solve the discrete nonlinear system of equations. Specifically, the exact Jacobian is formed at each iteration and this linear system is then solved iteratively with a preconditioned Krylov method for each Newton iteration. In this preliminary work we focus on improving the efficiency and scaling of this approach. We comment on Jacobian-Free Newton-Krylov methods as well as other preconditioners in our conclusions below.

The exact Jacobian offers a significant challenge for iterative solvers. In particular it may be written in the form,

$$\mathcal{J} = \mathcal{D} + \mathcal{N}$$

where \mathcal{D} is a symmetric block diagonal operator (the Picard linearization) and \mathcal{N} is nonsymmetric with no derivatives. The difficulty posed by this system is three-fold. First the coefficients of entries in \mathcal{N} involve gradients of the current target metric tensor, and hence, can be either positive or negative. In addition, their magnitude may be large, implying strong coupling of the components. Finally, the exact Jacobian is nonsymmetric.

Thus far, our work has focused on evaluation of solver performance in the 2D research code that was developed for [1]. This code used BPKIT to provide a preconditioner comprised of Block ILU (1) in conjunction with two global passes of SSOR for FGMRES. While performance of this solver is adequate for small test problems, it is clear that it will not scale well enough to handle the large 2D and 3D problems of interest to the ASC program. In particular, the linear scaling in solution cost required by large-scale simulations can only be provided by robust multilevel solver and preconditioning algorithms. Thus, the primary objective of our preliminary investigation was to demonstrate that the nonsymmetric linear system involving the full Jacobian could be preconditioned effectively with the well-known Ruge-Stüben serial Algebraic Multigrid Solver (AMG). To this end we have provided the option to use AMG both as a linear solver and as a preconditioner for GMRES in this research code. It is important to note that this is the basis for the parallel Los Alamos Algebraic Multigrid (LAMG) Solver by Wayne Joubert (LA-CC 03/107).

The preliminary performance study was conducted on a rectangular horseshoe domain with an initial grid defined by transfinite interpolation. This provided a logically rectangular grid of 64×64 points. A relative convergence criteria for the Newton iteration of 10^{-6} in the two norm is used for all solvers. The performance of AMG as a preconditioner is quite impressive, with a

bounded number of iterations per Newton step for both single V(1,1) and W(1,1)cycles, as well as for the more practical W(1,1) cycles, with a relative convergence criteria of 10^{-2} . This is clearly shown in the top figure, while the comparison with the BPKIT preconditioner is shown in the bottom figure. Most notably, for the BPKIT preconditioner the number of GMRES iterations/Newton iteration increases linearly as the Newton iteration converges. In addition to this poor scaling it is important to note that BPKIT preconditioned GMRES iterations will grow with increasing grid size, while the number of AMG preconditioned GMRES iterations should remain constant with increasing grid size.

This preliminary study has demonstrated the significant gains in efficiency that are possible by using advanced solvers such as AMG preconditioned GMRES. In particular, we have developed a solution algorithm that should scale linearly with the number of grid points. In the near future we hope to conduct more detailed scaling and robustness studies. In addition, future work will consider using the Jacobian-Free approach in conjunction with the use of the Picard linearization, \mathcal{D} , as a preconditioner. This combination has been shown to work very well for scalar quasilinear diffusion equations. Furthermore, this preconditioner is symmetric and block diagonal, and hence, it fully decouples into n scalar systems. Thus, solving the preconditioning system with AMG will be very efficient.

- [1] G. Hansen, A. Zardecki, D. Greening, and R. Bos, "A Finite Element Method for Unstructured Grid Smoothing," *J. Comp. Phys.* **194** (2), 611–631 (2004).
- [2] G. Hansen, A. Zardecki, D. Greening, and R. Bos, "A Finite Element Method for Three-Dimensional Unstructured Grid Smoothing," Los Alamos National Laboratory report LA-UR-03-8671 (2004).




Figure 1—
Smoothing is performed on a rectangular horseshoe domain with an initial grid defined by transfinite interpolation. In the top plot, three different cycling options are shown for AMG preconditioned GMRES, demonstrating that the number of GMRES iterations / Newton iteration is bounded. This is in contrast to the linear growth in the number of BPKIT preconditioned GMRES iterations, which are shown in the lower plot.